

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

| | | |
|-----------------------------|---|------------------------|
| In re: Application of: | : | |
| Brown et al. | : | |
| | : | Before the Examiner: |
| Serial No: 10/099,777 | : | Cam Linh T Nguyen |
| | : | |
| Filed: 03/14/2002 | : | Group Art Unit: 2171 |
| | : | |
| Title: APPARATUS AND METHOD | : | Confirmation No.: 4836 |
| OF EXPORTING FILE SYSTEMS | : | |
| WITHOUT FIRST MOUNTING THE | : | |
| FILE SYSTEMS | : | |

APPELLANTS' REPLY BRIEF UNDER 37 CFR §41.41

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This is a Reply to the Examiner's Answer dated November 13, 2007 in accordance with 37 CFR §41.41.

RESPONSE TO EXAMINER'S ARGUMENTS

In the Answer Brief of November 13, 2007, the Examiner asserted that the consulting step of the independent claims is taught by Vahalia et al. in col. 13, lines 19 – 26. Appellants respectfully disagree.

In col. 13, lines 19 – 39, Vahalia et al. disclose:

Each of the data movers in a cluster has a directory of the file systems in the cluster and a data base of the mount points for the file systems and the data mover owner of each read-write file system. When any data mover in a cluster receives a request for access, it checks the data base, and if it finds that the file system is a read/write file system having a different primary data mover, then the data mover forwards the client request to the primary data mover. Read requests, write requests, mount requests, and lock requests are examples of data access requests that are forwarded to the primary data mover for the file system to be accessed. If a data mover receives a request for access to a file system in the cluster and finds that the request is a read-only request or the data mover finds that it is the primary data mover for the file system, then it accesses the file system directly. Look-up of the file name in the file directory and management of access authorization for a file system can be distributed between a secondary data mover that receives and interprets a client request for access to the file system and the primary data mover that manages the locks on the file system.

Thus, in the above-reproduced paragraph, Vahalia et al. disclose that each computer system in a cluster of computer systems that can grant access to a file system to a requesting client has a directory (or file) that includes all file systems in the cluster and which computer system is allowed to grant access to which one of the file systems. This allows computer system₁ that receives a

request for a file in file system₂ to determine that access to file system₂ can only be granted by computer system₂ and to forward the request to computer system₂.

Thus, the directory or file is not associated with a mount point but rather is associated with an access-granting computer system in the cluster. Further, the directory or file is not consulted in order to export a file system but rather to determine which one of the computer systems in the cluster can grant access to the file system and thus to which computer system the request is to be forwarded.

Consequently in the cited paragraph, Vahalia et al. do not teach the step of ***consulting a file associated with a mount point of a mounted file system to retrieve needed information to export the file systems, the mount point being the point at which the file systems are mounted on a computer system*** as in the claimed invention.

The Examiner further asserted that Vahalia et al. teach in col. 13, lines 55 – 58 the step of exporting the file systems. Appellants again disagree.

In col. 13, lines 51 – 64, Vahalia et al. disclose:


In step 94, execution branches to step 95 if the file for which access is requested is in a remote file system. This occurs if a "mount point" is reached during file look-up as the file name, which is a path through the directory tree, is traversed from the root of the tree to the file to be accessed. In step 95, the data mover sends the request to an NFS thread for export to the remote file system corresponding to the "mount point." If the remote file system is in the cached disk storage subsystem (23 in FIG. 7), then the mount request is exported by forwarding it to the data mover that is the owner of the file system of the "mount point." If the process of file-name path traversal ends at the file to be accessed without ever reaching a "mount point," then execution continues from step 94 to step 96.

Thus in the above-reproduced paragraph, Vahalia et al. disclose that when a request is received by a computer system for access to a file, the computer system has to figure out in which one of the file systems in the cluster the requested file exists. This is done by traversing the pathname of the file from a root directory to the file to be accessed. If while traversing the pathname, a mount point is reached, then that means the file resides in a remote file system. Consequently, the computer system will use an NFS thread to send the request to the remote file system. If the remote file system has already been cached, a mount request is sent to the computer system that can grant access to the remote file system.

However, Vahalia et al. do not teach in that paragraph the step of **exporting file systems** as claimed.

Hence, Appellants submit that the teachings of Vahalia et al. do not anticipate the claimed invention and request reversal of the rejection.

Respectfully Submitted

By: 
Votel Emile
Attorney for Applicants
Registration No. 39,969
(512) 306-7969